Plaintiffs' Exhibit 110

Admeld Migration Design Document (Early Draft)

Status: Draft (as of 9/14/2012)
Author: jsolinsky@google.com (Jason W. Solinsky)

This is a work in progress. There is NOTHING in this document that is not subject to change as of the current date!

Objective

Rapidly migrate the Admeld publishers to AdX while minimizing the disruption to their businesses.

Background

Google acquired Admeld for roughy $400M in December 2011. The Admeld platform combines Real Time Bidding with numerous features that Google Ad Exchange did not previously support, including Traditional Yield Management, Private Exchanges, and Bid Landscapes. In the nine months since the acquisition, AdX has made substantial progress towards replicating these features.

In the first half of 2013 The Admeld migration team will migrate all or substantially all of the Admeld customers to AdX. A rapid migration will allow AdX engineering to maintain its focus on innovation and feature enhancements. At the same time, a smooth migration will minimize the number of customers that we lose during the migration process (although some customer loss is inevitable).

The key to both a speedy and smooth migration is automation. This document will evolve into the design document for the software used to automate the Admeld to AdX migration.

Schedule: Four Phases

The Admeld to AdX migration will occur in four phases:

Phase 1: Prototyping and Planning: During this phase we will demonstrate that all of the technologies needed to complete a migration are available (or identify and remedy any gaps). At the completion of this phase we will have three things:
1. A working prototype that is able to migrate test accounts from Admeld into AdX for the simplest cases.
2. A detailed design document outlining the work to be performed during phase 2. [*The document you are reading will evolve into the phase 2 design document.*]
3. A project plan, identifying every task that must be completed during phase 2, identifying who will perform it, and how much time it will require. This may be as simple as a well documented tracker instance, or it may be substantially more sophisticated. (TBD)

During Phase 1, we will determine whether or not the Admeld migration can be performed using AdX rules based inventory.

Although Phase 1 was originally scheduled to end during the first half of November, it will be our goal to complete it during the first half of October.

Phase 2: Development and Testing: During this phase we will develop and test a version of the migration tools that is robust enough to migrate Admeld's most important and finicky clients (NBC, Quadrant One, etc.) Although it is aggressive (especially since phase one is still ongoing), we will commit to being Phase 2 code complete in our Q4 OKRs.

Phase 3: Migration: During this phase we will migrate or terminate substantially all of Admeld's customers. Phase 3 will end on June 30, 2013.

Phase 4: Clean up: During this phase we will migrate any stragglers (it is expected that a handful of the largest publishers will manage to defer their migration past the end of Phase 3). We will then export the Admeld reporting data, and shut down the Admeld servers no later than December 31, 2013.

Scope

A great deal of Admeld data is not going to be migrated automatically (or at all). The data not being migrated includes:
- All Buyside customer data. Only Publishers are being migrated.
- All reporting data.
- Prioritized Bids/Private Exchanges. (TAMs will be responsible for reimplementing these in AdX where appropriate.)
- Automatic Floors
- Mobile tags
- AEM
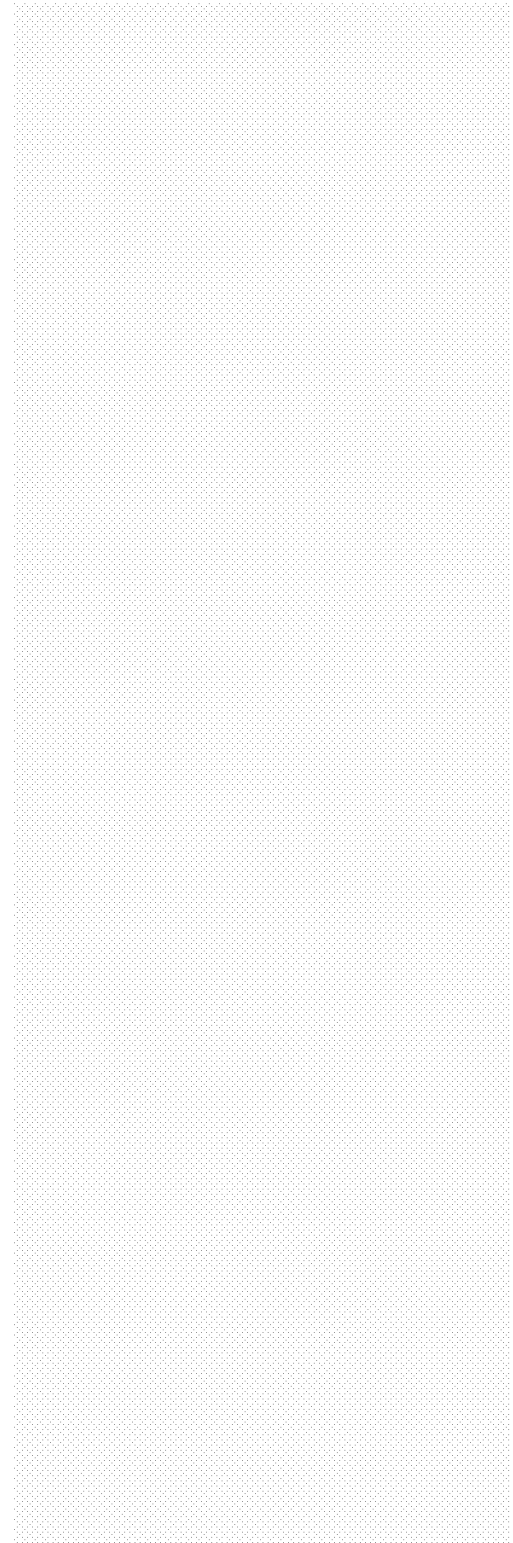
The automatic tools will migrate the following:
- Inventory information (Admeld websites and placements which will become AdX rules ofrTracking Code Channels and Ad Units
- CPM Floors
- Block Lists
- Network Tag information (Network Tags, Network Accounts and Scrape Script Parameters)
- User account data (Who is permitted to log into the system, with what permissions?)

Overview

The automated migration of each Admeld customer will involve the following steps:

1. Setup: An AdsDb AdX Customer, and AdX web property will be created (if none already exists) manually before automated setup can begin. Objects shared by many customers (such as BuyerNetworks) will also be created on the AdX and TYM FE systems.

2. Initiation: The Technical Manager of the migration will fill out a web form containing: the Admeld Customer Id, the AdX web property code, and a "migration configuration" (the latter selected from a short drop down menu). It is possible that additional information will be required for certain configuration choices. Once this form is submitted, the export will be initiated by sending a GET request through periscope.

3. Export: Code running on Admeld Servers (leveraging the existing process of generating JANIML for the Admeld Ad Servers) will export the publisher's data into protobufs that are returned to Google via HTTP.

4. Inventory and some RTB Import: The structure of the Publisher's Admeld inventory (Web Sites, Placements [and potentially CPM Thresholds and Block Lists]) will be analyzed and converted to Ad Units or Rules. If the publisher is an existing AdX publisher, it may be necessary to modify existing AdX inventory, but the default (in the absence of clear requirements to the contrary) will be a purely additive migration in which Admeld inventory is added alongside existing AdX inventory, and any newly created Rules to not overlap with existing inventory [**confirming this or identifying the manipulations required on existing AdX inventory is the highest priority uncertainty at this moment and could have large impacts on the migration process**]. A map from the Admeld Inventory to the newly created units and rules will be created for use during TYM migration. The ASFE/AdX Inventory Service will be used to create the new inventory.

5. Additional RTB import: CPM thresholds and Blocks not created during the previous step will be created using the Inventory Service and ARC Service.

6. TYM Import: In a single Megastore Job (all TYM FE objects are clustered together by Publisher) each publisher's NetworkTags and NetworkAccounts will be imported into TYM FE. The existing TYM FE DAO methods will be called directly by the import process. Adwords import, will be handled automatically by the TYM FE Reconciliation Process.

7. Inspection: Once these steps have been completed, TAMs and clients will review their new settings in AdX ICS and TYM FE and confirm that they are correct.

Commented [1]: How will we handle GAIA? Will the user have a "wizard" to walkthrough on first login? If so, we'll need to educate some publishers on this step.

Commented [2]: We don't know a task is in tracker. Because this is non-blocking, I would expect to put it off a few more weeks.

Commented [3]: We have, in fact, decided to deliberately defer addressing this until November at the earliest.

Commented [4]: +jsolinsky@google.com +jarid@google.com FYI some of the code used to migrate adx data from doubleclick. https://critique.corp.google.com/#review/22325712

8. Going Live: The Publisher's ad server will replace the existing Admeld Tags with AdX tags. For XFP customers, dynamic allocation to AdX _may_ be used as the default method. In some cases the priority of XFP dynamic allocation may have to be altered from the default (12) to achieve the same effect as is achieved from existing Admeld tags.

9. Celebration: Alcohol is imbibed and revenue increases.

10. (9b.) Fixes:

Detailed Design

Documentation

i18n
There will be no i18n of the migration process.

Latency

Scalability

Migrations will be performed one at a time. The largest customers will require significant amounts of memory on both sides of the process, but the requirements are small when considered from the perspective of Google scale.

Logging

We will log all exceptions during the import phase of the migration process, and examine them before we begin the inspection phase

Testing Plan

Work Estimates